arXiv:1412.5720v1 [cs.MS] 18 Dec 2014

FlexDM: Enabling robust and reliable parallel data mining using WEKA

Madison Flannery[1], David M Budden[2] and Alexandre Mendes[3]

[1] Melbourne Graduate School of Science, University of Melbourne, Parkville, VIC 3010, Australia
[2] Systems Biology Laboratory, University of Melbourne, Parkville, VIC 3010, Australia
[3] School of Electrical Engineering and Computer Science, The University of Newcastle, Callaghan, NSW 2308, Australia
Alexandre.Mendes@newcastle.edu.au

Abstract. Performing massive data mining experiments with multiple datasets and methods is a common task faced by most bioinformatics and computational biology laboratories. WEKA is a machine learning package designed to facilitate this task by providing tools that allow researchers to select from several classification methods and specific test strategies. Despite its popularity, the current WEKA environment for batch experiments, namely *Experimenter*, has four limitations that impact its usability: the selection of value ranges for methods options lacks flexibility and is not intuitive; there is no support for parallelisation when running large-scale data mining tasks; the XML schema is difficult to read, necessitating the use of the Experimenter's graphical user interface for generation and modification; and robustness is limited by the fact that results are not saved until the last test has concluded.

FlexDM implements an interface to WEKA to run batch processing tasks in a simple and intuitive way. In a short and easy-to-understand XML file, one can define hundreds of tests to be performed on several datasets. FlexDM also allows those tests to be executed asynchronously in parallel to take advantage of multi-core processors, significantly increasing usability and productivity. Results are saved incrementally for better robustness and reliability.

FlexDM is implemented in Java and runs on Windows, Linux and OSX. As we encourage other researchers to explore and adopt our software, FlexDM is made available as a pre-configured bootable reference environment. All code, supporting documentation and usage examples are also available for download at `http://sourceforge.net/projects/flexdm`.

1 Introduction

Large-scale data analysis is an integral component of bioinformatics and computational biology research. Genome-based studies (*e.g.* microarray and next generation sequencing) allow tens of thousands of gene expression levels to be quantified and analysed simultaneously ([1,2]). In such situations, sophisticated statistical and data mining tools have to be used to fully understand the massive amounts of data generated. WEKA (*Waikato Environment for Knowledge Analysis*) is one of those tools. Designed for machine learning and data mining, WEKA is a widely-used gold standard framework that has been cited more than 7,000 times in scientific literature ([3]).

When a specific dataset is to be analysed, it is difficult to know *a priori* which combination of classifier and hyperparameters will provide a better understanding of the data. Consequently, WEKA data mining experiments typically involve the comparison of multiple classifiers to identify which best approximates the unknown function relating input features (epitomised by the recent work of [4]). To facilitate this process, WEKA provides a module called *Experimenter*. Even though the Experimenter partially automates some data mining tasks, we have identified four major limitations with its functionality, as described in the following section.

2 Approach

FlexDM addresses the following major limitations in the default WEKA Experimenter:

- **Limitations to value ranges for methods options**: WEKA allows the setting of value ranges for hyper-parameters by using the meta-classifiers `CVParameterSelection` and `GridSearch`. However, they allow only one or two method options to be tested at a time, respectively. In addition, the setup of those meta-classifiers requires several nonintuitive steps. FlexDM imposes no such limitations; it allows any number of value ranges, and all combinations are tested in parallel.
- **Lack of parallel processing**: WEKA does not inherently utilise parallel processing when executing data mining tasks. Given that all mainstream CPUs are multicore with hyper-threading, this represents a missed opportunity for increased performance and productivity. Parallel processing for WEKA has previously been achieved with other packages(*e.g.* Weka-Parallel [5] and Grid-enabled Weka [6]) and has been incorporated as part of the FlexDM functionality accordingly, using asynchronous parallel processing to achieve maximum performance.
- **Reliance on the Experimenter GUI**: WEKA uses an XML schema to represent the set of experiments to be executed, which gives a high degree of flexibility and extensibility to the tool ([7,8]). However, the schema used by WEKA is very difficult to understand and modify without using the software's GUI. FlexDM uses a simplified XML schema with improved readability, allowing users to rapidly change test specifications without the need for a GUI.
- **Lack of robustness**: The Experimenter runs all tests in sequence, but results are saved only after the last one has concluded. In real world applications, such tests often take days to weeks of processing, which means that a hardware failure will cause the loss of all intermediate results. FlexDM saves results incrementally as each job concludes, allowing the user to restart from the last task successfully completed.

3 Example Application

The FlexDM XML schema is intuitive and easily modified. Figure 1 shows an example input file for the execution of two classifiers (*J48* and *PART*), with one hyperparameter (*confidence factor* - C) varying from 0.1 to 1.0 in increments of 0.1. This input file specifies leave-one-out cross-validation on the dataset "*health.arff*". The confusion matrix will be saved in each results file. FlexDM also creates a *summary file* containing the classification accuracy for each test, allowing easy identification of the best performing classifier.

The compact representation of FlexDM is enabled by the assumption that any information not specified (*e.g.* hyperparameters) defaults to WEKA's default values. The file in Figure 1 is 11 lines long and has 321 characters (excluding spaces), 10x more compact than the WEKA equivalent (available on the supplementary material webpage).

```
<!DOCTYPE flexdm SYSTEM "flexdm.dtd">
<flexdm>
   <dataset name="health.arff" test="leavexval" results = "matrix">
      <classifier name="weka.classifiers.trees.J48">
         <parameter name="-C" value="[0.1:0.1:1.0]" />
      </classifier>
      <classifier name="weka.classifiers.rules.PART">
         <parameter name="-C" value="[0.1:0.1:1.0]" />
      </classifier>
   </dataset>
</flexdm>
```

Fig. 1. Example of XML input file for FlexDM. The schema used is intuitive and easily modified. The option "C" is given a range [0.1:0.1:1.0], resulting in 10 independent tests for each classifier with C = 0.1, 0.2, ..., 1.0. Comparatively, the equivalent WEKA XML file has 131 lines and 6,367 characters.

Another feature of FlexDM is its asynchronous parallel processing. If an input XML file specifies more than one classifier, the software runs them in parallel using exactly (n-1) logical core units present in the host

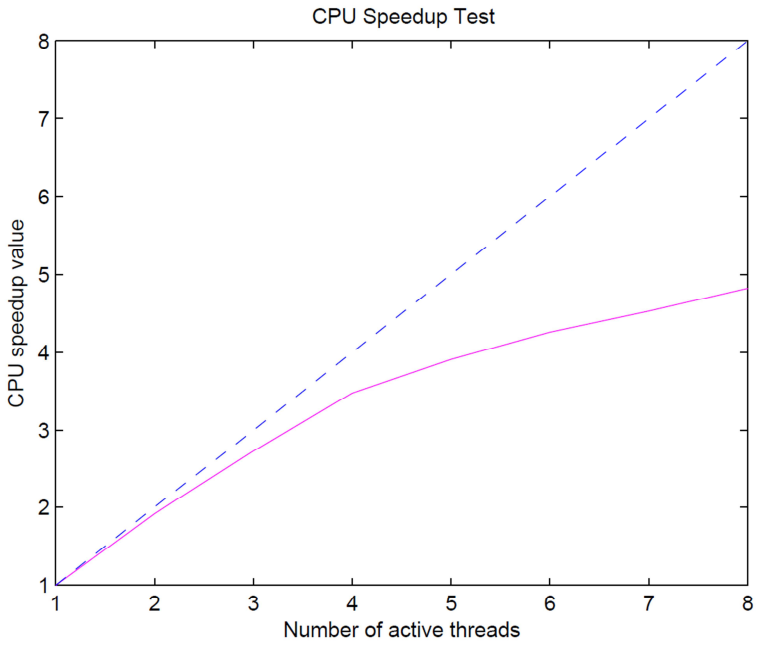

Fig. 2. Speedup generated by FlexDM's asynchronous parallel processing (pink) compared to theoretical maximum (blue), using a quad-core CPU.

computer (one core is left available for other activities). In Figure 2, we show the speedup obtained with the XML test file from Figure 1. The test was carried out on an Intel i7 processor (quad-core CPU with hyper-threading). Note the quasi-linear speedup up until four threads, followed by a less steep speedup curve once hyper-threading is used.

4 Implementation

As we encourage other researchers to explore and adopt our software, FlexDM is implemented using open-source software and made available as a pre-configured bootable virtual environment using the approach described by [9]. This environment was created using a minimal installation of Lubuntu 13.10; a lightweight Linux distribution which supports all the tools required. OpenJDK version 2.4.7 and WEKA version 3.6.10 were installed, along with the core set of packages and utilities required to reproduce our usage examples. Full configuration details for the reference environment are available in the Supplementary Materials. Alternatively, all data and scripts are available online at `http://sourceforge.net/projects/flexdm`.

5 Conclusion

This paper presents a software tool for robust and reliable parallel data mining using WEKA, FlexDM, which addresses four major limitations of the default WEKA Experimenter. The use of a simplified XML schema also makes FlexDM far easier to learn than the Experimenter. This tool can significantly improve productivity for any research group that frequently performs large-scale data mining or machine learning tasks, such as those commonplace in the bioinformatics and computational biology domain.

References

1. Budden, D.M., Hurley, D.G., Crampin, E.J.: Predictive modelling of gene expression from transcriptional regulatory elements. Briefings in Bioinformatics **doi: 10.1093/bib/bbu034** (2014)
2. Hurley, D.G., Cursons, J., Wang, Y.K., Budden, D.M., Crampin, E.J., et al.: NAIL, a software toolset for inferring, analyzing and visualizing regulatory networks. Bioinformatics **doi: 10.1093/bioinformatics/btu612** (2014)
3. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA data mining software: an update. ACM SIGKDD Explorations Newsletter **11**(1) (2009) 10–18
4. Marsden, J., Budden, D., Craig, H., Moscato, P.: Language individuation and marker words: Shakespeare and his Maxwell's Demon. PloS one **8**(6) (2013) e66813
5. Celis, S., Musicant, D.R.: Weka-parallel: machine learning in parallel. In: Carleton College, CS TR, Citeseer (2002)
6. Khoussainov, R., Zuo, X., Kushmerick, N.: Grid-enabled Weka: A toolkit for machine learning on the grid. ERCIM news **59** (2004) 47–48
7. Achard, F., Vaysseix, G., Barillot, E.: XML, bioinformatics and data integration. Bioinformatics **17**(2) (2001) 115–125
8. Varde, A., Rundensteiner, E., Fahrenholz, S.: XML based markup languages for specific domains. In: Web-based Support Systems. Volume 2. Springer (2010) 215–238
9. Hurley, D.G., Budden, D.M., Crampin, E.J.: Virtual reference environments: A simple way to make research reproducible. Briefings in Bioinformatics **doi: 10.1093/bib/bbu043** (2014)